

Bayesian identification of protein differential expression in multi-group isobaric labelled mass spectrometry data

Howsun Jow, Richard J. Boys and Darren J. Wilkinson*

July 25, 2014

Abstract

In this paper we develop a Bayesian statistical inference approach to the unified analysis of isobaric labelled MS/MS proteomic data across multiple experiments. An explicit probabilistic model of the log-intensity of the isobaric labels' reporter ions across multiple pre-defined groups and experiments is developed. This is then used to develop a full Bayesian statistical methodology for the identification of differentially expressed proteins, with respect to a control group, across multiple groups and experiments. This methodology is implemented and then evaluated on simulated data and on two model experimental datasets (for which the differentially expressed proteins are known) that use a TMT labelling protocol.

Keywords: proteomics, Bayesian, mass spectrometry, MS/MS, differential protein expression

1 Introduction

In recent years, isobaric molecular labelling techniques have been developed which, in conjunction with tandem mass spectrometry (also known as MS/MS), can be used to perform quantitative analyses of complex protein mixtures (Thompson et al., 2003; Ross et al., 2004). This procedure involves using enzymes such as trypsin to digest the proteins in the samples into their constituent peptides. Isobaric molecular

*School of Mathematics & Statistics, Newcastle University, U.K.

labels (one label per sample) are then chemically attached to the resulting peptide components of the enzyme digested proteins. The samples are then mixed and passed through a liquid chromatograph into a mass spectrometer where tandem mass spectrometry is performed. The relative intensities of these labels for the resulting MS/MS spectra can be determined from the intensities of the isobaric labels' "reporter" ions. The MS/MS spectra are those of the constituent peptides and these can be identified either by hand or, more usually, using a protein sequence database. If the identified peptides are prototyped for a protein, i.e. peptides that can be uniquely derived from the digestion of a protein, then the relative amounts of this protein in the labelled samples can be quantified.

The experimental process described above is inherently stochastic as even technically identical replicates will produce different data regardless of the accuracy of the mass spectrometer. Also there is a wide "dynamic" range in the intensity of the detected peptides. This matters because only a limited number of ions can be subject to MS/MS. The highest intensity ions are the ones which tend to be selected. As such, the more abundant proteins, especially large proteins, will tend to drown out less abundant smaller proteins. A variety of techniques exist by which the complex protein mixture can be split up into simpler fractions but these do not completely eliminate the problem. An additional problem is that there are a limited number of isobaric tags available. The maximum currently commercially available is eight (Choe et al., 2007). This limits the number of samples that can be analyzed in a single experiment. However, it should be possible to analyze a larger number of samples using multiple experiments that, for example, share a common reference sample. This requires that any analysis technique used be able to link data from multiple experiments of this type.

Apart from the peptide identification problem (Befekadu et al., 2009), the main statistical problem in analysing MS/MS data is how to detect which of the proteins in a complex mixture are present in significantly different amounts in a set of isobaric labelled samples from two or more pre-determined groups such as a control group and a set of treatment groups. Currently standard software such as MASCOT (Perkins et al., 1999) is quite limited as they attempt to detect differences between two groups either by a informal thresholding on the fold-ratio or by using a t-test, which in turn can be generalised to a comparison of more groups via a one-way ANOVA. The t-test method typically makes the reasonable assumption that the intensities of the reporter ions for MS/MS spectra of proteotypic peptides are log-normally distributed (Boehm et al., 2007). However, using the t-test approach has several drawbacks. The main problem arises from having to test a large number of hypotheses to determine significant differences in expression levels between groups, as a test is needed for each of the hundreds of proteins. Of course, such problems can be somewhat mitigated by using standard multiple hypotheses

corrections (Sidak, 1968; Holm, 1979; Benjamini and Hochberg, 1995). Another problem with using t-test is that many suffer from low power due to proteins being quantified on the basis of only a small number of MS/MS spectra.

To circumvent these issues, various authors have developed methods using a more detailed model for the peptide intensity and use ANOVA analytic techniques to determine differences between protein expression levels for different groups and their associated statistical significance. Keshamouni et al. (2006) describe a very simple ANOVA model for normalised log-ratios in a single experiment comparing a single treatment against a control. A more sophisticated ANOVA model for multiple treatments and experiments is described in Oberg et al. (2008). Although their full model is somewhat overparameterised, it is easily reduced to an identifiable model by combining some parameters and setting others to zero. This reduced model is very closely related to the model described in this paper. The main innovations introduced here are that the model is fit jointly to all of the data simultaneously, rather than using a stepwise regression approach, that the fitting is done in a fully Bayesian way, and that the Bayesian version of the model includes variable selection indicators for differential expression, allowing direct inference to be made for the probability of differential expression associated with each protein. Note that related methods for LC-MS data are often tailored to specific features of that data. For example, the models of Karpievitch et al. (2009) and Wang et al. (2012) are tailored to a specific kind of informative missingness in the data which does not arise for isobaric labelled MS/MS data. In the MS/MS framework considered here, missingness is generally much less of an issue than for LC-MS data (Oberg and Mahoney, 2012), and there is little evidence to suggest that missing values in MS/MS data sets are strongly informative. Nevertheless, as for all proteomics technologies, there is still an issue with certain peptides not being detected at all during the first phase of the MS/MS procedure. Our model deals correctly with such missingness, in a fairly efficient way.

In this paper we fit a detailed random effects model with an ANOVA form to analyse protein expression levels, but remove or consolidate many of the parameters that Oberg et al. (2008) find to be typically confounded. The model is a hierarchical model which borrows strength appropriately across multiple peptides, proteins, samples and experiments. Unlike the stepwise regression approach described by Oberg et al. (2008), our fitting approach does not force each protein to have an independent variance parameter, and correctly propagates normalisation uncertainty without assuming approximate orthogonality of model components. In common with Oberg et al. (2008), our model allows for the analysis of multiple experiments with common reference samples. Additionally, our method adopts a fully Bayesian approach and therefore has all the advantages of interpretability and being able to include prior information (where available). The Bayesian approach also has sev-

eral advantages specific to the context of complex hierarchical models. The framework is flexible, and allows convenient checking for over-parametrisation or variable confounding by prior to posterior comparisons and multi-dimensional analysis of the posterior distribution. Importantly the model has a variable selection form which ensures that we fit an appropriate model for the combination of differentially expressed and non-differentially expressed proteins. Models without this structure have the drawback that they inflate the error variance in the “null” model due to contamination by outlying differentially expressed proteins which then hinders the detection of differential expression.

In the next section we give more details of the experimental framework and describe the model together with the prior-to-posterior analysis. In section 3 we demonstrate the potential of the model to pool information across multiple MS/MS experiments by using simulated data. In sections 4 and 5 we analyse two real MS/MS datasets. The real datasets have a simple structure but serve to demonstrate that our model captures the behaviour of real experimental data. The dataset in section 4 has two technically identical replicates and so contains a negative control, that is, there should be no differentially expressed proteins. The dataset in section 5 is provided by the ProteoRed consortium as part of an assessment of various quantitative proteomics methods. It contains known differentially expressed proteins within an otherwise set of technically identical replicates. The paper concludes in section 6 with a discussion.

2 Methods

2.1 Experimental framework

The experimental framework of an isobaric labelled MS/MS experiment is as follows: the protein samples, each containing roughly equal amounts of total protein by mass, are digested using enzymes with typically high specificity such as trypsin i.e. they cleave the protein molecules into peptide fragments at predictable amino acid residues. The digested samples are then each chemically labelled with different isobaric labelling molecules. The samples are then mixed together and run through a liquid chromatograph into a mass spectrometer.

A selection of high intensity ions from the resulting MS spectra are then shattered, using high energy collisions, to yield further ion spectra called MS/MS spectra. The chemically attached molecules leave signature or reporter ions in the MS/MS spectra. The relative intensity of these reporter ions will give us relative quantitative information about the amounts of the peptide fragments, corresponding to each MS/MS spectra, in each sample.

The above constitutes a single experiment. Multiple experiments (E) can of course be performed either with technical or biological sample replicates. The samples being analyzed are typically assigned to pre-determined biological groups. For simplicity, only the MS/MS spectra corresponding to peptide fragments uniquely derived from a single protein are used in the quantitative analysis. This means that a single MS/MS spectrum gives quantitation information about a single protein. Thus we need a model that will simultaneously model the observed expression levels for P proteins for the G biological groups to which the samples being analyzed are assigned. One of the biological groups is used as the control group. Differences in protein expression level are then measured relative to this control. The total protein content of each sample that ends up being labelled are only roughly equal. In our model we include a parameter which allows for differing amounts of total labelled samples in an experiment. In addition, when performing the analysis across multiple experiments it is important to be able to “link” the experiments. In order to facilitate this a reference sample is selected in each experiment. Ideally this should be a technical replicate of some standard sample that is included in all our experiments.

An example of an experimental design is given in Figure 1. It shows a scenario with two experiments for twelve samples from four groups with three replicate samples per group assuming that only six isobaric labels are available for any single experiment.

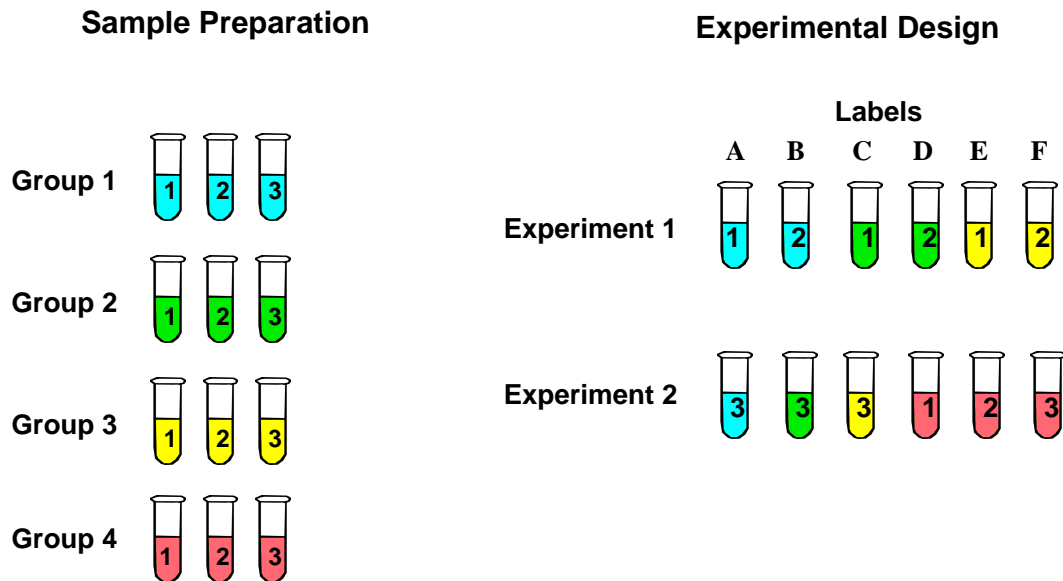


Figure 1: Schematic of an experimental design for twelve samples from four groups (three replicate samples per group) using six available isobaric labels (A-F).

2.2 The Model

We consider a series of isobaric labelled mass spectrometric experiments designed using n_I distinctly identifiable isobaric tags (we will typically have $n_I = 6$). We assume that for each experiment e , we have (log) intensity measurements y_{egjki} in treatment group g arranged by sample/replicate number i from an MS/MS spectrum k identified as being that of a peptide derived from a protein j . Specific peptides are therefore identified by a particular (j, k) combination. The isobaric tags are used to identify the (g, i) combinations within an experiment. The model has an ANOVA style, starting with a sample-specific normalising constant, κ_{egi} , and then decomposing the remaining uncertainty using a variable selection approach in the style of Kuo and Mallick (1998); see O’Hara and Sillanpaa (2009) for a review of other possible approaches. The benefits of Bayesian hierarchical modelling in the context of gene expression microarrays is already well established (Hein et al., 2005). Explicit inclusion of the normalisation constant in the model ensures that important information in the data is not lost during the normalisation process (Hein et al., 2005; Karpievitch et al., 2012; Callister et al., 2006; Hill et al., 2008). The benefits of including the normalising constant directly in ANOVA-style fixed and random effects models for proteomics data is most clearly articulated in Oberg et al. (2008), to which the reader is referred for further details.

We model the log intensities as

$$y_{egjki} = \kappa_{egi} + \alpha_{jk} + \beta_{gj}\gamma_{gj} + \epsilon_{egjki}, \quad (1)$$

$e = 1, \dots, E$, $g = 1, \dots, G$, $j = 1, \dots, P$, $k = 1, \dots, m_j$ and $i = 1, \dots, n_{eg}$, where, for identifiability, the parameters are subject to the constraints

$$\beta_{1j} = \gamma_{1j} = 0, \quad j = 1, \dots, P \quad \text{and} \quad \kappa_{eg_e 1} = 0, \quad e = 1, \dots, E,$$

where g_e is the group in experiment e whose first sample is used as the reference sample. In general, differential expression will be measured against a control group and this group will be labelled group 1. The control group will commonly also provide a reference sample in each experiment (giving $g_e = 1$), though sometimes the experimenter might choose different groups to provide the reference samples. Thus the number of isobaric labels used in each experiment is $\sum_g n_{eg} = n_I$. Note that a MS/MS spectrum is assumed to be “assigned” to one and only one protein.

In a typical experiment some reporter ions do not show up in a particular MS/MS spectrum. This corresponds to missing data in the model and can be dealt with by treating these unobserved values in the usual Bayesian MCMC way, that is, by including the missing data in the model as unobserved parameters and stochastically imputing them along with all other unknowns. This leads to the correct posterior distribution (which is less concentrated than it would have been in the case

of full observation), but at the expense of poorer mixing of the MCMC scheme. However, a more efficient implementation is possible by using a “ragged array” with varying dimension indices, which avoids creating variables corresponding to the missing data. We adopt this latter scheme in the implementation we describe in section 2.3.2, as it has significantly better mixing properties than the naive approach.

The interpretation of the parameters is as follows. The parameters κ_{egi} are normalization constants for sample i in group g of experiment e relative to the reference sample in that experiment, that is, the log-ratio of the total amount of sample i (in group g of experiment e) with respect to the reference sample (sample 1 in group g_e of experiment e). The α_{jk} parameters are the mean log-intensities of a peptide k (corresponding to a reporter ion in an MS/MS spectrum) assigned to protein j for labelled samples from (control) group 1. The β_{gj} parameters are binary indicating whether or not the protein j is differentially expressed for biological group g with respect to (control) group 1 and the γ_{gj} parameters are the difference in the mean log-expression level of protein j between group g and (control) group 1. Finally, the ϵ_{egjki} are independent and identically distributed noise terms following a normal $N(0, \sigma^2)$ distribution with zero mean and variance σ^2 .

2.3 Bayesian Inference

The inference task is to make statistically valid statements about the unknown model parameters $(\kappa, \alpha, \beta, \gamma, \sigma)$ that describe the unknown normalisation factors, the mean expression levels in the different groups and the experimental noise. The Bayesian statistical inference approach combines information from the data $D = \{y_{egjki}\}$ with that from prior information using Bayes Theorem, and describes this through the posterior distribution. If we assume that the prior distribution for each model parameter can be specified independently, the posterior distribution is

$$\begin{aligned} \pi(\kappa, \alpha, \beta, \gamma, \sigma | D) \propto & \prod_{e, (g, i) \neq (g_e, 1)} \pi(\kappa_{egi}) \times \prod_{jk} \pi(\alpha_{jk}) \times \prod_{g \neq 1, j} \pi(\beta_{gj}) \pi(\gamma_{gj}) \times \pi(\sigma) \\ & \times \sigma^{-n} \exp \left\{ -\frac{1}{2\sigma^2} \sum_{egjki} (y_{egjki} - \kappa_{egi} - \alpha_{jk} - \beta_{gj}\gamma_{gj})^2 \right\}, \end{aligned}$$

where $n = E n_I \sum_j m_j$ is the total number of log-intensity measurements for the isobaric labels and $\pi(\cdot)$ denote prior probability (density) functions.

2.3.1 Prior Distribution on Model Parameters

The Bayesian approach allows for additional information to be incorporated by using a prior distribution on all model parameters $(\kappa, \alpha, \beta, \gamma, \sigma)$. Although each analyst should incorporate their own prior beliefs into the prior distribution, we have chosen to represent prior beliefs by using standard distributions. This leaves the analyst to decide on their own choice for the parameters of these distributions. The prior distributions are, for $e = 1, \dots, E$, $g = 1, \dots, G$, $j = 1, \dots, P$, $k = 1, \dots, m_j$ and $i = 1, \dots, n_{eg}$

$$\begin{aligned}\kappa_{egi} &\sim \text{N}(a_\kappa, 1/b_\kappa), & \alpha_{jk} &\sim \text{N}(a_\alpha, 1/b_\alpha), \\ \beta_{gj}|p_{gj} &\sim \text{Bern}(p_{gj}), & p_{gj} &\sim \text{Beta}(a_p, b_p), \\ \gamma_{gj} &\sim \text{N}(a_\gamma, 1/b_\gamma), & \sigma^{-2} &\sim \text{Ga}(a_\sigma, b_\sigma),\end{aligned}$$

where $\text{Ga}(a, b)$ is a gamma distribution with mean a/b , $\text{Bern}(p)$ is a Bernoulli distribution with success probability p and $\text{Beta}(a, b)$ is a Beta distribution with mean $a/(a + b)$. The essential model structure over the unknowns is displayed as a plate diagram in Figure 2.

We suggest the following default choices for the prior parameters. In an ideal series of experiments, the experimenter uses the same total amount of reporter in each experiment. In other words, experimenters try to design their experiments so that the κ_{egi} are close to zero. We input this information by taking $a_\kappa = 0$ and $b_\kappa = 1/9$. In our experience, log-intensities for reporter ions typically range between 4 and 16 and so taking this range as four standard deviations gives $a_\alpha = 10$ and $b_\alpha = 1/9$. In most analyses of this kind, the proportion of differentially expressed proteins will be low, say around 5%. It is possible to undertake the analysis with this proportion fixed at such a user-specified value. However, as results may be very sensitive to the value chosen, we prefer to put a prior distribution on the proportion and learn about it from the data. We suggest taking $a_p = 1$ and $b_p = 19$.

Naturally the degree of differential expression in proteins (relative to the control group) will vary but we anticipate a typical fold-change to be around one (zero on the log-scale) and so recommend taking $a_\gamma = 0$ and $b_\gamma = 1$. The level of measurement accuracy of the reporter ion log-intensity measurements is particularly difficult to assess. Therefore we suggest taking a quite weak prior distribution for σ by using $a_\sigma = b_\sigma = 1/1000$. Overall we regard these choices of prior parameters as inputting relatively weak prior information. Some analysts may feel they have stronger views to include in their prior distribution and so may feel justified in taking quite different choices to these “default” ones. We will use our default choices in the subsequent analyses we present in this paper.

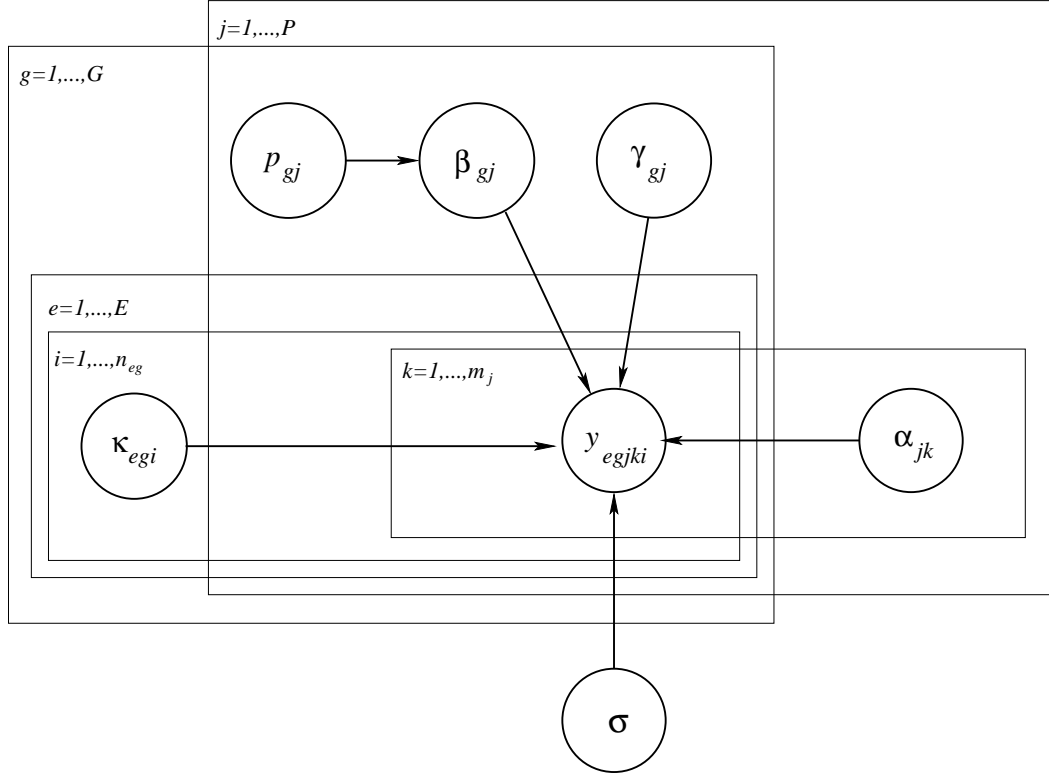


Figure 2: Plate diagram for the essential model structure. The directed acyclic graph (DAG) over the random variables represents the factorisation of the joint distribution of unobserved and observed model variables, and is used to construct posterior inference schemes. Fixed hyperparameters are omitted for clarity.

2.3.2 Posterior analysis

Unfortunately, the posterior distribution is analytically intractable for any significant number of proteins and peptides. However, it is possible to simulate realizations from this distribution by using computer-intensive Markov chain Monte Carlo (MCMC) techniques; see Gamerman and Lopes (2006). Essentially this algorithm simulates a Markov chain which has the posterior distribution as its equilibrium distribution. Thus, after the algorithm has converged (in distribution), all subsequent realizations will have the required (posterior) distribution. We have chosen to implement our model using the free open source software JAGS (Plummer, 2003). The JAGS code for our model and the key steps in this Gibbs sampling scheme are given in the appendix. Note that the indexing used in this code is slightly different to that in (1) to deal with missing data more efficiently: it uses a “ragged array”

structure to avoid creating nodes for missing values. We have found that this fairly simple one-variable-at-a-time Gibbs updating scheme works very effectively for this class of models and so more sophisticated MCMC schemes have not been pursued. We have also developed a full R package implementing the model and this is available as the R-Forge package **dpeaqms** (dpeaqms, 2011). In more challenging scenarios, single variable updating schemes for variable selection problems can be inefficient, and some authors have proposed methods (such as collapsing and then joint updating) which can improve mixing in this case (Bottolo et al., 2010; Davies et al., 2014). However, such methods are not integrated into off-the-shelf MCMC engines such as JAGS.

3 Simulated Data Analysis

In order to evaluate the performance of the statistical inference technique outlined previously, the methods were tested on data simulated from the model (1). The experimental scenario we consider is one in which there are $G = 4$ groups consisting of a control group (CTL) and three treatment groups (TRT1, TRT2 and TRT3) and we have $n_I = 6$ isobaric tags to use. We assume that differential expression is determined with respect to the control group and so label this as group $g = 1$; the other groups TRT1, TRT2 and TRT3 we label as $g = 2, 3, 4$ respectively. The simulated dataset we constructed consists of 12 labelled samples and, as this must have resulted from more than one experiment, we assume that there were $E = 2$ experiments, with six tagged samples per experiment. We also assumed that overall there were three samples from each group. The samples were assigned to the six tag reporter ions according to the pattern (CTL, CTL, TRT1, TRT1, TRT2, TRT2) for the first experiment and (CTL, TRT1, TRT2, TRT3, TRT3, TRT3) for the second experiment. Therefore, the numbers of samples per group in each experiment are $n_{11} = n_{12} = n_{13} = 2$, $n_{14} = 0$ and $n_{21} = n_{22} = n_{23} = 1$, $n_{24} = 3$. Note that this experimental design is not as balanced as it might have been, for example, we could have required at least one sample from each group in each experiment. However it might not always be possible to construct “balanced” experiments due to having a large number of groups relative to the number of tags available, and so we investigate here the effect of using an “unbalanced” experimental design. As we have a sample from (control) group 1 in each experiment, we take these samples as the reference samples for the experiments, that is, take $g_1 = g_2 = 1$.

We constructed the dataset to contain the results on $P = 300$ proteins, with the number of simulated MS/MS spectra per protein (m_j) drawn from a geometric distribution with mean 6, this distribution being chosen to be roughly comparable with the real MS/MS datasets analyzed in this paper. The mean log-intensities of

the reporter ions for (control) group 1 (α_{jk}) were drawn from a normal distribution with mean 10 and standard deviation 3.0. Again, this distribution is similar to those for the real datasets in the paper. The probability of a protein being differentially expressed with respect to the control group CTL was drawn from Bernoulli distributions with probabilities 0.1, 0.2 and 0.3 for treatment groups TRT1, TRT2 and TRT3 respectively and resulted in 26, 71 and 98 differentially expressed proteins in the treatment groups. We considered fold changes in differential expression relative to the (control) group varying between 1.5 and 4.0 for up-regulation and between 1/4.0 and 1/1.5 for down-regulation. This led to us drawing the parameters for the level of differential expression (γ_{gj}) from a uniform distribution on $(-1.39, -0.41) \cup (0.41, 1.39)$. The log ratios of the amount of protein in each sample with respect to a reference sample for each experiment (sample 1 in group g_e of experiment e), the κ_{egi} , were drawn from a normal distribution with mean 0 and standard deviation 0.1. Finally the variability of the errors (σ) was set to 0.3.

The data are summarized in Figure 3, with the left-hand column giving summaries for experiment 1 and the right-hand column for experiment 2. Figures 3(a) and (b) show histograms of the number of MS/MS spectra per protein for the two experiments. Clearly the number of proteins decreases as the number of spectra per protein increases, and is consistent with the simulation model in which the number of spectra per protein follows a geometric distribution with mean 6. Figures 3(c) and (d) are box plots of the log-intensities of the reporter ions (labelled by group and sample within group) in the two experiments. The plots show that there are no clear differences between groups/samples, with these levels dominated by the overall mean levels (α_{jk}), which in this simulation have mean 10 and standard deviation 3. Figures 3(e) and (f) are MA-plots for two selected samples (sample 1 in group 2 (TRT1) of experiment 1 and sample 1 in group 4 (TRT3) of experiment 2). These plots display differences between these samples and the reference sample (in the corresponding experiment) in the log-intensities (m) against their mean (a) for each MS/MS spectra in an attempt to highlight any dependence between variability and overall level in the data. These plots suggest (correctly) that there is no such dependence in these data.

3.1 Results

We analysed these data using the prior distribution with the default parameter options and looked at the output of the JAGS code using a variety of initial starting points. Typically, 10-30K iterations were needed to attain distributional convergence (burn-in) but for the analyses in this paper we adopt a conservative strategy and use a burn-in of 100K iterations. Note that convergence was assessed by using

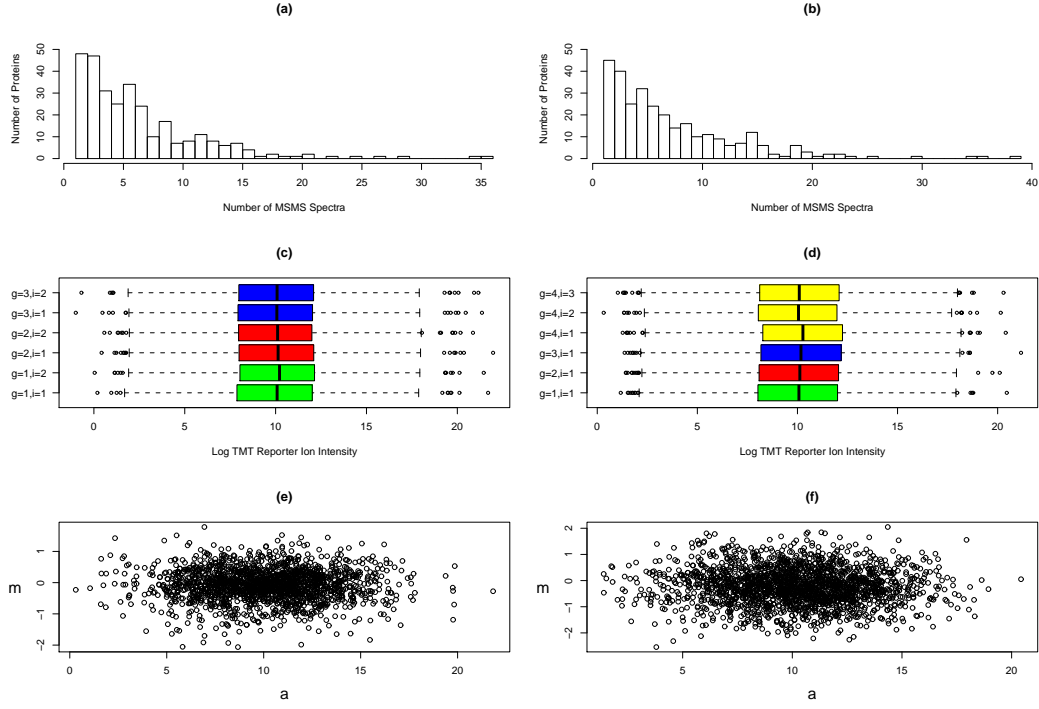


Figure 3: Summary plots for the simulated data: left panel – experiment 1, right panel – experiment 2. (a) and (b) Histograms of the number of MS/MS spectra per protein; (c) and (d) Box-plots of the reporter ion log-intensities by group/sample; (e) and (f) MA-plots for two selected samples (showing the difference in log-intensities, m , against the average log-intensity, a).

a variety of informal and formal tests. We report here the results from five independent runs of the JAGS code in which (after burn-in) each chain was run for a further 100K iterations and thinned by taking every 100th iterate. This gives a total of 5K (weakly correlated) realizations from the posterior distribution for analysis.

Figure 4 shows kernel density plots for an MCMC sample from the posterior probability distribution for a representative selection of parameters: the measurement error standard deviation σ , four of the normalisation constants κ_{egi} and a log-fold difference in intensity with respect to (control) group 1 $\beta_{gj}\gamma_{gj}$ for a protein which was differentially expressed in this group. The vertical dashed lines show the value used in generating the simulated data. Their location on the posterior plots verifies that, despite inputting fairly weak prior information, the posterior analysis has recovered these values reasonably accurately. Note that there is no “spike” at zero in the posterior density of $\beta_{2,116}\gamma_{2,116}$ because there were no zero values of $\beta_{2,116}$ in the posterior sample.

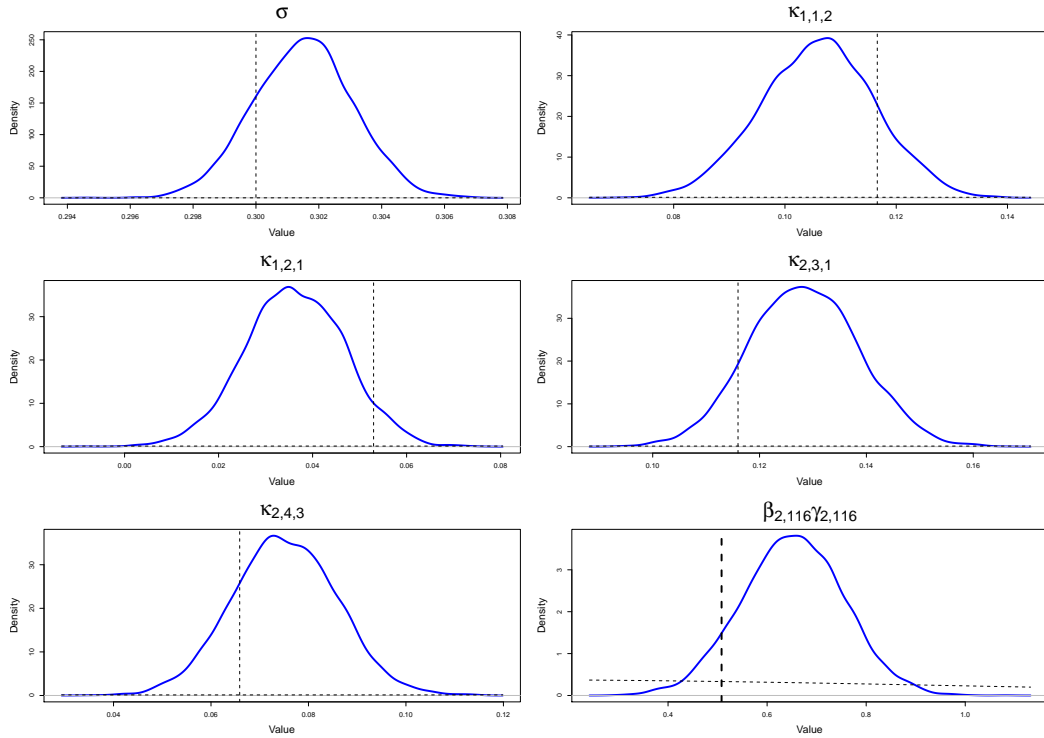


Figure 4: Posterior kernel density plots for a selection of parameters: the measurement error standard deviation σ , four of the normalisation constants κ_{egi} and a log-fold difference in intensity with respect to (control) group 1 $\beta_{gj}\gamma_{gj}$ for a protein which was differentially expressed in this group. The vertical dashed lines indicate the value used in generating the simulated data and the dashed curves are the prior densities used in the analysis.

Table 1 summarises the overall performance of the inference method in identifying the differentially expressed proteins. Classification is based on a posterior probability of differential expression exceeding 0.5, as is common practice for Bayesian variable selection problems. In fact the classification is rather insensitive to the particular choice of probability threshold, since the vast majority of proteins have posterior probability of differential expression either less than 0.1 or greater than 0.9. The full distribution of posterior probability of differential expression for the proteins in each of the 3 treatment groups is summarised in Figure 5. Despite the relatively high noise level, the method performs quite well for TRT1 and TRT2 but less well for TRT3. On further inspection, we found that the method generally failed to correctly identify a differentially expressed protein when either the level

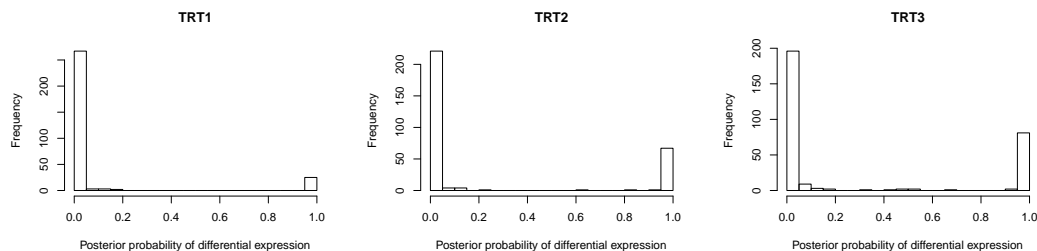


Figure 5: Histograms showing the posterior probability of differential expression for the proteins in groups TRT1, TRT2 and TRT3, respectively. The vast majority of proteins in the 3 samples have posterior probability very close to 0 or 1.

Group TRT1		Simulated with DE	Simulated without DE	
Inferred as with DE		24	1	25
Inferred as without DE		2	273	275
		26	274	300

Group TRT2		Simulated with DE	Simulated without DE	
Inferred as with DE		70	0	70
Inferred as without DE		1	229	230
		71	229	300

Group TRT3		Simulated with DE	Simulated without DE	
Inferred as with DE		87	0	87
Inferred as without DE		11	202	213
		98	202	300

Table 1: Performance of the method in detecting differential expression (DE) in proteins between groups and the control group. Inference for DE is based on a classification threshold of a posterior probability exceeding 0.5.

of differential expression was small or there were only one or two MS/MS spectra for the protein in each experiment. Additional complications for TRT3 were due to its samples only being present in the second experiment.

4 Case Study 1: Dataset on human plasma

Our first case study analyses data from an MS/MS analysis of human plasma published in Dayon et al. (2010). The data arise from an experiment to investigate the use of a particular mass spectrometric technique in the identification and quantification of peptides. The experimental design consists of a single experiment ($E = 1$) and produced two technically identical samples. We consider these samples to be single samples ($n_{11} = n_{12} = 1$) from two “artificial” groups ($G = 2$). Thus, any differentially expressed proteins will be due to variability in the experimental process and the technique used to detect differential expression, and so we expect to find relatively few proteins that are differentially expressed between the two groups.

The experiment was conducted as follows. After reduction, alkylation and digestion of the sample with trypsin, two technically identical sub-samples were taken, labelled with TMT-2plex labels (i.e. $n_I = 2$) and mixed. The resulting mixture was then run through a liquid chromatograph into a Tandem Mass Spectrometer and MS/MS spectra acquired. In this simple experiment, it does not really matter which group supplies the reference sample. Here, we take this as coming from group 1, that is, $g_1 = 1$. The MS/MS spectra, in the form of Mascot Generic Format (MGF) files, were kindly made available for this study by Alexander Scherl. These were then analysed using Proteome Discoverer version 1.1 (Thermo Fisher Scientific) and version 3.65 of the IPI sequence database (Kersey et al., 2004). The list of identified MS/MS spectra was then filtered by requiring them to pass three criteria: (i) the peptide needed a high identification confidence, with the threshold calculated to give a false discovery rate on the decoy database of 0.01, (ii) the MS/MS spectra needed to be identified as being those of peptides uniquely derivable from a single protein, and (iii) the MS/MS spectra needed to include both reporter ion peaks. This left 3158 MS/MS spectra, corresponding to peptides from $P = 94$ proteins, for analysis.

Figure 6 gives some summary plots of the data. Figure 6(a) gives a histogram of the number of MS/MS spectra per protein. Clearly there are many proteins with only a few spectra and only a few with many spectra. Note that this shape is consistent with the geometric distribution assumed in the simulation study. Figure 6(b) gives a box plot of the log-intensities of the reporter ions (labelled by group and sample within group) and shows no obvious differences between the groups. Figure 6(c) gives the MA-plot for the two samples. The plot is similar to those in the simulation study (Figures 3(e) and (f)). Here there is a slight suggestion of increased variability associated with small measurements, but the effect seems to be small, and so for the purposes of this analysis we will continue to assume constant variability.

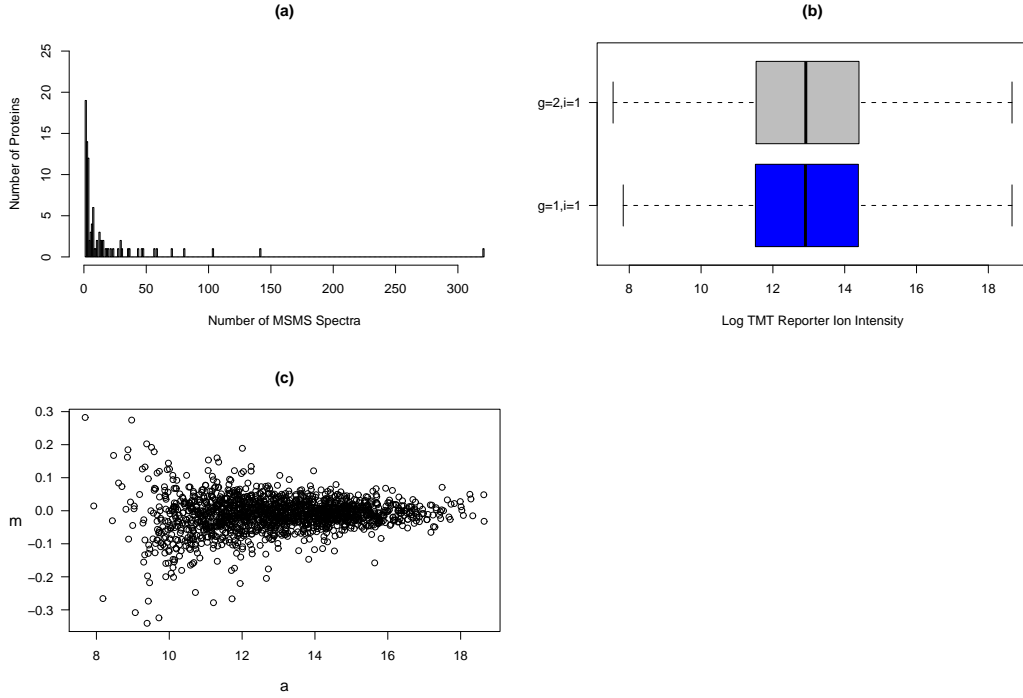


Figure 6: Summary plots of the Dayon et al. (2010) human plasma data. (a) Histogram of the number of MS/MS spectra per protein; (b) Box-plot of the TMT 2-Plex reporter ion log-intensities by group/sample; (c) MA-plot of the log-intensities for the two samples (showing the difference in log-intensities, m , against the average log-intensity, a).

4.1 Results

We now analyse these duplex TMT reporter ion log-intensities and report here the results from runs of the JAGS code using the procedure described in section 3.1 which give 5K (weakly correlated) realizations from the posterior distribution.

Figure 7 shows summary plots for an MCMC sample from the posterior probability distribution for a representative selection of parameters: the measurement error standard deviation σ , the (only) normalisation constant $\kappa_{1,2,1}$ for the second sample ($g = 2, i = 1$) with respect the reference sample ($g = 1, i = 1$) and the log-fold difference in intensity $\beta_{2,62}\gamma_{2,62}$ for protein IPI00647704 with respect to (control) group 1. This protein was chosen as it has the highest posterior probability of being differentially expressed. The trace and auto-correlation plots are typical of all five chains and suggest that there are no mixing problems and that convergence has been attained.

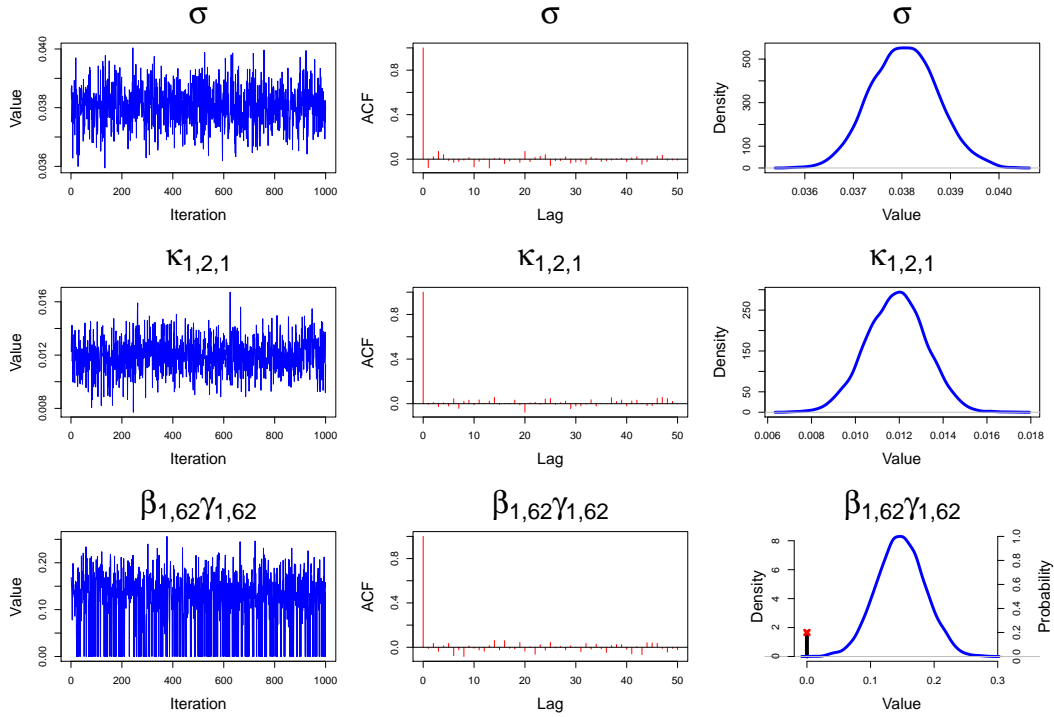


Figure 7: Results for an analysis of the data in Dayon et al. (2010): trace and auto-correlation plots from a single chain and the posterior kernel density plot from all five chains for a representative selection of parameters: the measurement error standard deviation σ , the log-normalisation parameter $\kappa_{1,2,1}$ and the product $\beta_{2,62}\gamma_{2,62}$ for protein IPI00647704.

Means of the posterior output for the differential expression indicator parameters β_{gj} can be used to estimate the posterior probability of differential expression for each protein. Of the 94 proteins examined, only two had a posterior probability exceeding 0.1. These were IPI00647704 and IPI00916434, with posterior probabilities 0.8 and 0.3 respectively. If we assume that the two samples were perfect technical replicates, then using a posterior probability threshold of 0.5 to declare whether or not proteins are differentially expressed gives a false positive rate around 1%.

A natural part of any data analysis is to assess the validity of the model used to make inferences. We favour checks using the posterior predictive distribution of the (logged) intensities, that is, their distribution allowing for the posterior un-

certainty in the model parameters; see Gelman et al. (2003), Chapter 6 for details. The posterior predictive density is straightforward to determine using the MCMC output $\{\kappa_{egi}^{(\ell)}, \alpha_{jk}^{(\ell)}, \beta_{gj}^{(\ell)}, \gamma_{gj}^{(\ell)}, \sigma^{(\ell)}; \ell = 1, \dots, N\}$ for our model as

$$f(y_{egjki}) \simeq \frac{1}{N} \sum_{\ell=1}^N \phi \left(\frac{y_{egjki} - (\kappa_{egi}^{(\ell)} + \alpha_{jk}^{(\ell)} + \beta_{gj}^{(\ell)} \gamma_{gj}^{(\ell)})}{\sigma^{(\ell)}} \right).$$

A useful diagnostic of model fit can be based on the location of the observed intensities within their individual predictive distributions as a well fitting model will produce posterior predictive distributions consistent with the observed log-intensities. Figure 8 shows histograms of samples from the posterior predictive densities for a random selection of nine observed log-intensities. These show that the model provides a good fit to the data (shown in red).

It is interesting to compare our results with those obtained by using existing methods: a t-test and a moderated t-test (Smyth, 2004), both using log intensity data, as in our Bayesian analysis. Since the t-test looks at differences between peptide intensities (on the log scale), the normalisation constant for specific peptides (corresponding to α in the Bayesian model) drops out of the analysis, and therefore does not require explicit normalisation. However, the normalisation constant associated with the sample (corresponding to κ in the Bayesian model) does not drop out, and therefore the raw log intensities must be pre-processed in order to apply sample normalisation before the t-test can be carried out. Here we normalise each sample by its mean value. This is one of the standard normalisation techniques commonly used, and is the method most directly comparable to our Bayesian model. Including the normalisation constants explicitly in the model has the advantage that it allows direct modelling of the experimental data. Our approach to modelling normalisation constants is similar to the approach used by Oberg et al. (2008) in a frequentist context. For these data, both the t-test and the moderated t-test suggest that there are no differentially expressed proteins at the 5% significance level after using an FDR correction for testing multiple hypothesis (Benjamini and Hochberg, 1995). It is worth noting that the pre-normalisation step is necessary, and that different methods of normalisation can lead to different results (Karpievitch et al., 2012). By contrast, the Bayesian model-based method developed here includes normalization constants explicitly as part of an integrated model and so normalizing the data prior to analysis is not required. We have also found that our method is relatively insensitive to whether or not the data have been preprocessed using standard normalization methods.

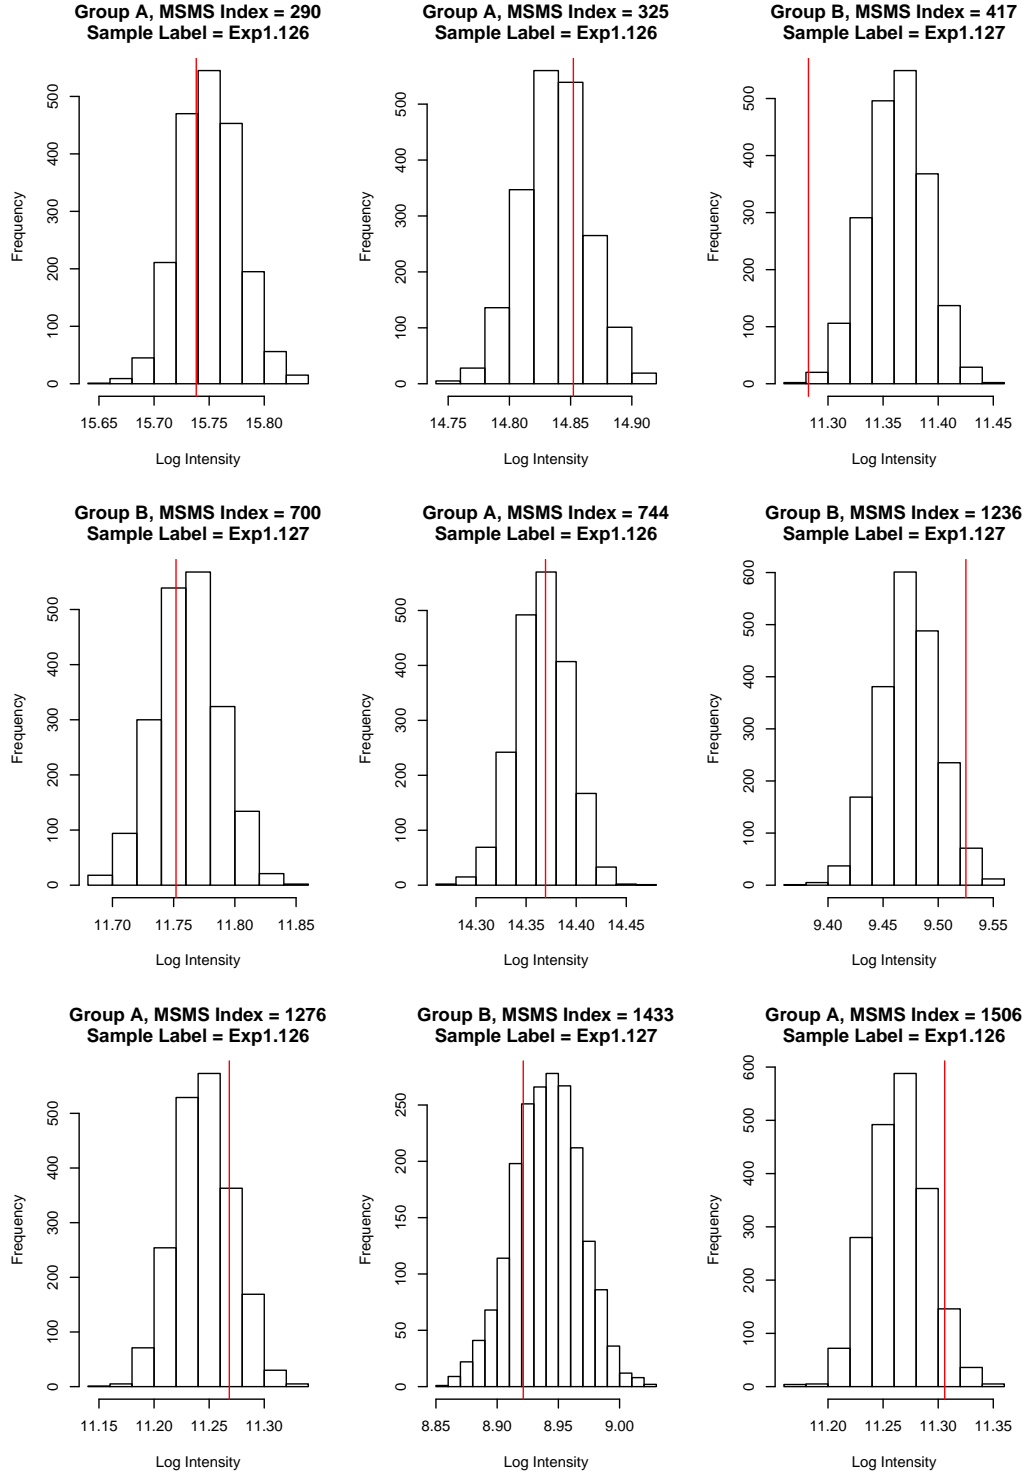


Figure 8: Histograms showing the posterior predictive distributions for a random selection of observed log-intensities (uniquely identifiable by experiment e , protein j , MSMS spectra k and sample gi). The red vertical line shows the observed log-intensity.

5 Case Study 2: ProteoRed multi-centric experiment 5

Our second case study analyses data from an experiment carried out by the ProteoRed consortium as part of an assessment of various quantitative proteomics methods. The experiment analysed a mixture of *E. coli* proteins. The mixture was prepared by fractionation of the cytoplasmic proteome of *E. coli* and contained soluble proteins with a wide range of values for the isoelectric point (pI, the pH at which a molecule carries no net electrical charge) and average molar mass (Mw).

After dividing the mixture into two identical portions A and B, four xenobiotic proteins were added into the portions in differing amounts. The aim of this research was to evaluate the performance of different proteomics methods by observing their ability to detect the existence of the xenobiotic proteins within the otherwise identical samples/portions. The xenobiotic proteins used were CYC_HORSE (Cytochrome C, Mw 12362), MYG_HORSE (Apomyoglobin, Mw 16952), ALDOA_RABIT (Aldolase, Mw 39212) and ALBU_BOVIN (Serum albumin, Mw 66430) and their (theoretical) differential expression between the two portions is given in Table 2.

Code	Protein	Ratio (B/A)	$\log(B/A)$
CYC	Horse Cytochrome C	1.50	0.4055
MYG	Horse Apomyoglobin	0.38	-0.9676
ALDOA	Rabbit Aldolase	0.50	-0.6931
ALBU	Bovine Serum Albumin	5.00	1.6094

Table 2: Theoretical ratios of xenobiotic proteins in portions A and B

The experiment produced data by using a TMT-6plex ($n_I = 6$) labelling of three technical replicate sub-samples of each of the two portions. Sub-samples in portion A were labelled with even TMT-6plex labels (TMT126, TMT128, TMT130) and, in portion B, labelled with odd TMT-6plex labels (TMT127, TMT129, TMT131). The resulting labelled sub-samples were then mixed and the mixture divided into two technically identical portions. Each of these portions was then subjected to independent MS/MS analyses. In terms of our model, this gives an experimental design consisting of $E = 2$ experiments and $G = 2$ groups, and we label portion A as group 1 and portion B as group 2. The numbers of replicates in each experiment \times group are $n_{11} = n_{12} = n_{21} = n_{22} = 3$. The reference sample was set to its default in both experiments ($g_1 = g_2 = 1$).

We obtained the MS/MS peak list as an MGF file for both of the analyses. These were then analysed using Proteome Discoverer v.1.1 (Thermo Fisher Sci-

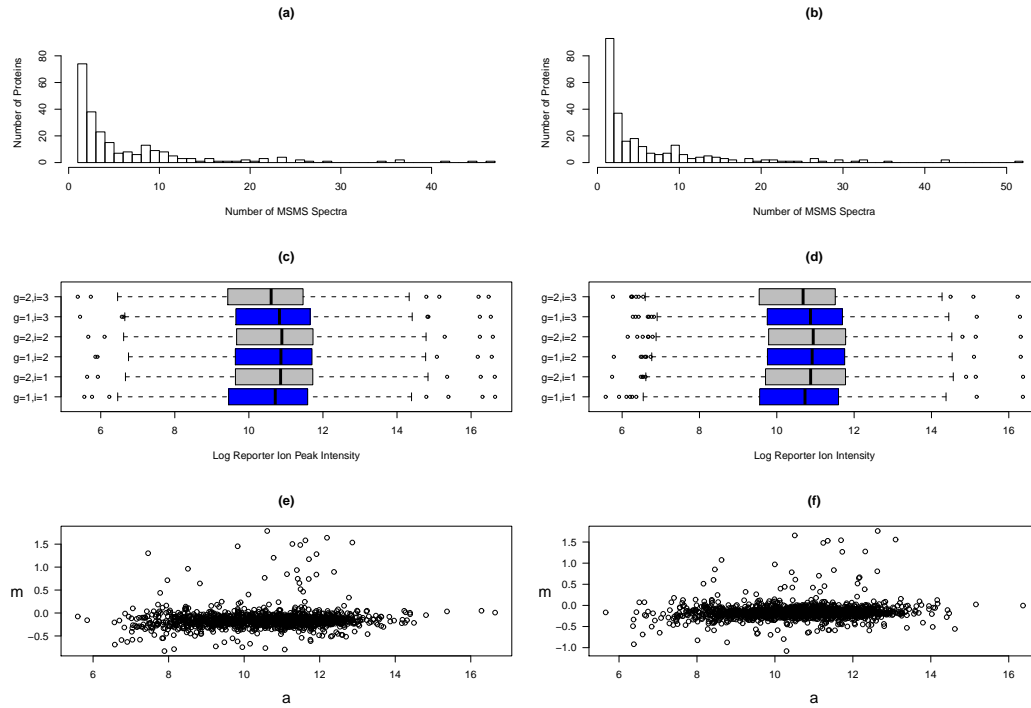


Figure 9: Summary plots for the ProteoRed data: left panel – experiment 1, right panel – experiment 2. (a) and (b) Histograms of the number of MS/MS spectra per protein; (c) and (d) Box-plots of the reporter ion log-intensities by group/sample; (e) and (f) MA-plots for two selected samples (showing the difference in log-intensities, m , against the average log-intensity, a).

entific) and the current version of the UniProt sequence database for *E. coli* (Jain et al., 2009). The list of identified peptides was then filtered to retain only those peptides with a high identification confidence, with the threshold calculated to give a false discovery rate on the decoy database of 0.01, and which were proteotypic. Also MS/MS spectra which did not have complete quantitative information for the six labels were excluded from the analysis. This left 238 and 259 proteins, each with at least one fully quantified MS/MS spectra, for the analyses of portions A and B respectively. In total, the two analyses provided quantitative information for 282 proteins. The data are summarised in Figures 9 and 10. Figures 9(a) and (b) shows histograms of the number of MS/MS spectra per protein for each experiment. As with the first case study, these distributions are consistent a geometric distribution. Figures 9(c) and (d) are box plots of the log-intensities of the reporter ions (labelled by group and sample within group), and panels (e) and (f) show MA plots for two of the samples. All of the plots suggest that the data is consistent with the

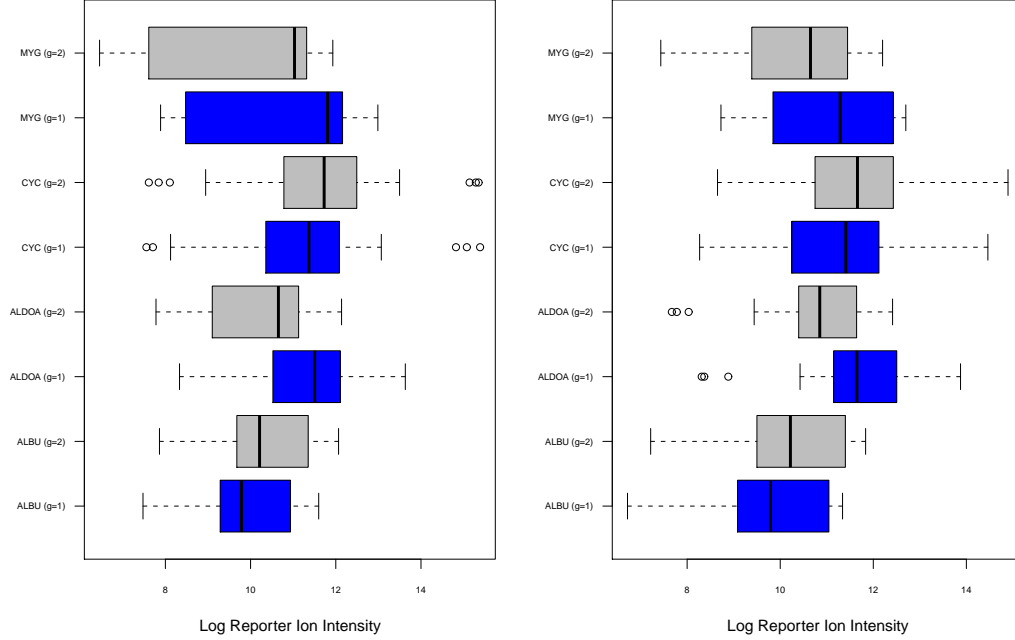


Figure 10: Box-plots of the reporter ion log-intensities by group for the four spiked in proteins: left panel – experiment 1, right panel – experiment 2.

assumptions of our model.

5.1 Results

We now analyse the ProteoRed dataset and, as with the previous analyses, base our analysis on the 5K realizations from the posterior distribution obtained by using the procedure described in section 3.1.

The analysis clearly identified all four spiked-in proteins as being differentially expressed, each with a posterior probability exceeding 0.999. The posterior distribution of the log ratio $\log(B/A)$ for each spiked protein is shown in Figure 11. Comparing these distributions with the theoretical values given in Table 2, we see that the values of the distributions for three of the four spiked proteins are of the right order of magnitude. Unfortunately, the analysis gives a posterior ratio for the ALBU (Bovine Serum Albumin) spiked protein which is out by a factor of more than 3. Other analyses of this dataset have come to similar conclusions (see results section in ProteoRed (2010)) and we suspect this is due to contamination in the

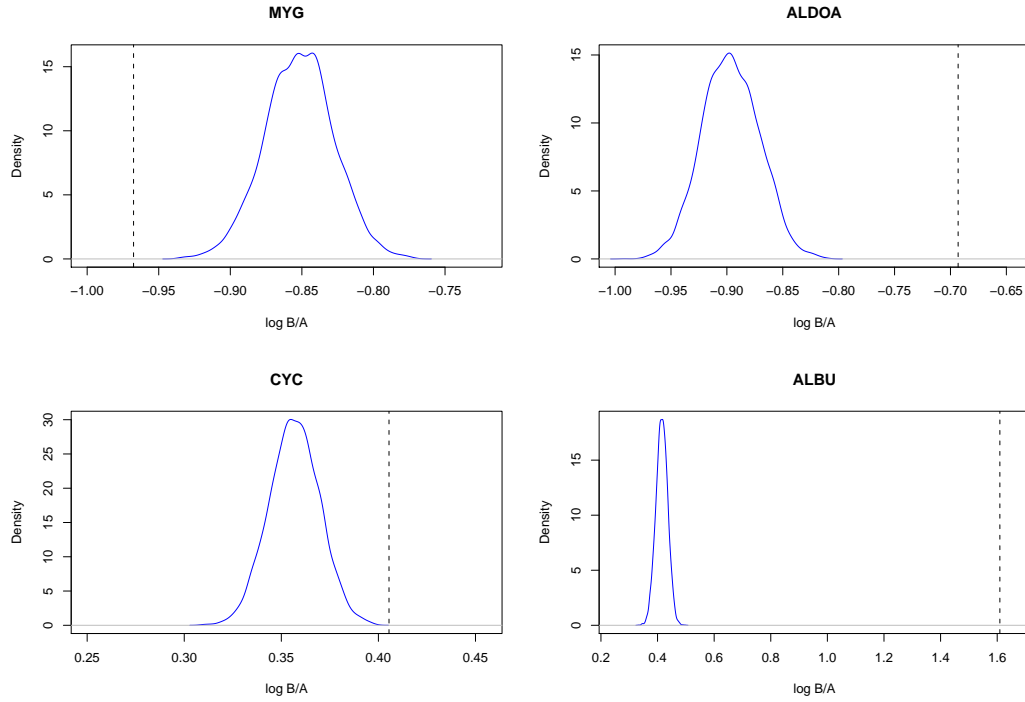


Figure 11: Posterior kernel density plots for the log ratio $\log(B/A)$ of the spiked proteins MYG, ALDOA, CYC and ALBU. The vertical dashed line shows the theoretical value for the log ratio.

original sample or at some stage in the preparation of the samples for the MS/MS experiment. Whatever the reason, this has clearly affected the estimation of the ratios for all the spiked proteins.

The analysis also identified four other (non-spiked) *E. coli* proteins as being differentially expressed when using a posterior probability threshold of 0.5: three main ones A7J0Q2, C3SFD7 and P02754 with posterior probabilities exceeding 0.999 and C3SVU2 with posterior probability 0.52. Nearly all of the remaining *E. coli* proteins (271 of the remaining 274) had a posterior probability of being differentially expressed below 0.05. As in the analysis of the dataset in section 4, we assessed model fit by comparing the observed data with their posterior predictive distributions and this confirmed that the model provided a reasonable fit to the data.

As before, we now compare the results of our analysis with those determined by using a standard t-test and by using a moderated t-test together with an FDR multiple hypotheses correction. Again, we follow standard practice and pre-normalise the data so that the sample means are the same. For both test-based analyses, the spiked proteins were found to be significantly differentially expressed

($q \leq 10^{-6}$). These analyses also found false positives, with the t-test analysis and the moderated t-test analysis identifying an additional ten and seven non-spiked proteins respectively ($q \leq 0.05$).

6 Discussion

In this paper we developed a model based approach to solving a general problem in quantitative proteomics using mass spectrometric (MS/MS) methods with isobaric labelling. The problem is that of determining which, if any, of a set of proteins in a complex mixture are differentially expressed between two or more groups and if so to quantify the degree of differential expression. This is not straightforward to resolve given the inherently stochastic nature of mass spectrometric data. Further complications are introduced, given that the protein intensities are not measured directly. What is observed are the intensities of the isobaric labels for the peptides resulting from enzyme digestion (usually trypsin) of these proteins.

The model based approach allows us to integrate data from multiple MS/MS experiments into a single statistical inference framework. An important feature of our framework is that these experiments do not have to be simple replicate experiments. This allows us to design a set of isobaric labellings that expands the effective number of samples (and therefore biological groups) that can be analysed beyond the maximum of eight isobaric labels which is currently commercially available (Choe et al., 2007). This ability to integrate data is a distinct advantage not offered by more standard approaches. Previous work by Hill et al. (2008) also uses a model based approach and makes this very point about combining data from multiple experiments. However, we use a more flexible and, we believe, more intuitively interpretable Bayesian statistical approach to inferring model parameters, whereas they use an ANOVA-based approach. An additional advantage of the Bayesian approach is that we can make direct inferences for the probability of differential expression, and so there is no need to correct for multiple hypotheses, as would be required for the ANOVA approach, for example, in order to determine whether protein ratios differed significantly from one.

A common concern with Bayesian analyses can be that results are sensitive to the choice of prior distribution parameters. In order to assess whether our results were robust and not an artifact of the prior choice, we investigated using a wide range of different values of the parameters of the priors. We found that the results for differentially expressed proteins were very robust for different prior choices for τ and κ . The results are slightly more sensitive to the prior choices for the distribution on p_{gj} , the prior probability of protein j in group g being differentially

expressed. Sensitivity to the prior on variable inclusion is a well-known issue in Bayesian variable selection (O’Hara and Sillanpaa, 2009).

In order to evaluate the method developed here we apply the technique to two MS/MS datasets. First the performance of the method was evaluated on a negative control. This is a dataset which we know should have no differentially expressed proteins. The dataset was from an MS/MS analysis of two technical replicate samples of human plasma labelled with a 2-plex TMT reagent. For this dataset we can reasonably assume that none of the proteins are differentially expressed. Of the 94 proteins identified from the data only a single protein was inferred as having a ≥ 0.5 mean posterior probability of being differentially expressed with a value of 0.815. This corresponds to an acceptably low false positive rate.

The second dataset was data from an MS/MS analysis of a sample used in a ProteoRed group experiment. The data consisted of technical replicates of an *E. coli* cytoplasmic proteome sample (A and B). These were spiked with differing amounts of four xenobiotic proteins. This acted as our positive control since the spiked proteins and their relative amounts are known. Again the method developed here was able to infer that the spiked proteins were differentially expressed with a high probability. The degree differential expression for these proteins was also of the right order for three of the proteins but was off by a factor > 3.0 for the Bovine Serum Albumin (ALBU) protein. Four out of the 278 non-spiked proteins were also inferred as differentially expressed: three with a high probability (> 0.999) and one with a borderline probability of 0.52.

In conclusion, our Bayesian statistical inference approach to determining differentially expressed proteins from isobaric labelled MS/MS data has been found to perform well on a variety of real and simulated datasets. The modelling framework allows us to perform a unified statistical analysis of multiple experiments and the multiple hypotheses are integrated into the model itself so there is no need for *ad hoc* normalisation methods or multiple hypothesis corrections. The model allows comparison of multiple experiments within a single unified model, thereby extending the range of applicability of the isobaric labelling technology. Importantly the model has a variable selection form which ensures that we fit an appropriate model for the combination of differentially expressed and non-differentially expressed proteins. Models without this structure have the drawback that they inflate the error variance in the “null” model due to contamination by outlying differentially expressed proteins which then hinders the detection of differential expression. A unified analysis approach using a frequentist ANOVA analysis has previously been described by Oberg et al. (2008). Fortunately, missing data is less of an issue when analysing isobaric labelled samples in MS/MS experiments than it is for other proteomic technologies, such as LC–MS. Nevertheless, missing values are often present in data sets, and this can complicate frequentist analysis consid-

erably. However, adopting a Bayesian approach is helpful in that it allows us to marginalise the model over any missing data as a routine part of the analysis. It would be relatively straightforward to extend our model to cover informative missingness, but we believe that our current model captures the most important sources of variation in typical labelled MS/MS datasets. Finally, the results we obtain are intuitively interpretable through simple probabilities of the analysed proteins being differentially expressed.

Acknowledgements

This work was funded by the UK Biotechnology and Biological Sciences Research Council through grants BBF0235451 and BBC0082001. The authors would like to thank Alexander Scherl (Proteomics Core Facility and Biomedical Proteomics Research Group, University of Geneva) for providing much of the data analysed in this paper, and to Satomi Miwa, Achim Treumann and Thomas von Zglinicki (Institute for Ageing and Health, Newcastle University) for helpful discussions. We would also like to thank the reviewers who have helped to improve the clarity of the paper.

References

- Befekadu, G. K., M. G. Tadesse, and H. W. Ressom (2009): “A Bayesian based functional mixed-effects model for analysis of LC–MS data,” in *Engineering in Medicine and Biology Society, 2009. EMBC 2009. Annual International Conference of the IEEE*, IEEE, 6743–6746.
- Benjamini, Y. and Y. Hochberg (1995): “Controlling the false discovery rate: A practical and powerful approach to multiple testing,” *Journal of the Royal Statistical Society, Series B*, 57, 289–300.
- Boehm, A. M., S. Putz, D. Altenhofer, A. Sickmann, and M. Falk (2007): “Precise protein quantification based on peptide quantification using iTRAQ,” *BMC Bioinformatics*, 8: 214.
- Bottolo, L., S. Richardson, et al. (2010): “Evolutionary stochastic search for Bayesian model exploration,” *Bayesian Analysis*, 5, 583–618.
- Callister, S., R. Barry, J. Adkins, E. Johnson, W. Qian, B. Webb-Robertson, R. Smith, and M. Lipton (2006): “Normalization approaches for removing systematic biases associated with mass spectrometry and label-free proteomics,” *Journal of Proteome Research*, 5, 277–286.
- Choe, L., M. D’Ascenzo, N. R. Relkin, D. Pappin, P. Ross, B. Williamson, S. Guertin, P. Pribil, and L. K. H. (2007): “8-plex quantitation of changes in

- cerebrospinal fluid protein expression in subjects undergoing intravenous immunoglobulin treatment for Alzheimer's disease." *Proteomics*, 7, 3651–3660.
- Davies, V., R. Reeve, W. Harvey, F. Maree, and D. Husmeier (2014): "Sparse bayesian variable selection for the identification of antigenic variability in the foot-and-mouth disease virus," in *Journal of Machine Learning Research: Workshop and Conference Proceedings*, volume 33, Journal of Machine Learning Research, volume 33, 149–158.
- Dayon, L., C. Pasquarello, C. Hoogland, J. C. Sanchez, and A. Scherl (2010): "Combining low- and high-energy tandem mass spectra for optimized peptide quantification with isobaric tags," *Journal of Proteomics*, 73, 769–777.
- dpeaqms (2011): "Bayesian statistical analysis of differential protein expression using data from MS/MS experiments on isobaric labeled samples," <http://r-forge.r-project.org/projects/dpeaqms/>.
- Gamerman, D. and H. F. Lopes (2006): *Markov Chain Monte Carlo: Stochastic Simulation for Bayesian Inference*, Chapman and Hall, 2nd edition.
- Gelman, A., J. Carlin, H. Stern, and D. Rubin (2003): *Bayesian Data Analysis*, Chapman & Hall/CRC, 2nd edition.
- Hein, A.-M. K., S. Richardson, H. C. Causton, G. K. Ambler, and P. J. Green (2005): "BGX: a fully Bayesian integrated approach to the analysis of Affymetrix GeneChip data," *Biostatistics*, 6, 349–373.
- Hill, E. G., J. H. Schwake, S. Comte-Walters, E. H. Slate, A. L. Oberg, J. E. Eckel-Passow, T. M. Therneau, and K. L. Schey (2008): "A statistical model for iTRAQ data analysis," *Journal of Proteome Research*, 7, 3091–3101.
- Holm, S. (1979): "A simple sequentially rejective Bonferroni test procedure," *Scandinavian Journal of Statistics*, 6, 65–70.
- Jain, E., A. Bairoch, S. Duvaud, I. Phan, N. Redaschi, B. Suzek, M. Martin, P. McGarvey, and E. Gasteiger (2009): "Infrastructure for the life sciences: design and implementation of the UniProt website," *BMC Bioinformatics*, 10: 136.
- Karpievitch, Y., A. Dabney, and R. Smith (2012): "Normalization and missing value imputation for label-free LC-MS analysis," *BMC Bioinformatics*, 13, S5.
- Karpievitch, Y., J. Stanley, T. Taverner, J. Huang, J. N. Adkins, C. Ansong, F. Hefron, T. O. Metz, W.-J. Qian, H. Yoon, et al. (2009): "A statistical framework for protein quantitation in bottom-up MS-based proteomics," *Bioinformatics*, 25, 2028–2034.
- Kersey, P. J., J. Duarte, A. Williams, Y. Karavidopoulou, E. Birney, and R. Apweiler (2004): "The International Protein Index: An integrated database for proteomics experiments." *Proteomics*, 4, 1985–1988.
- Keshamouni, V. G., G. Michailidis, C. S. Grasso, S. Anthwal, J. R. Strahler, A. Walker, D. A. Arenberg, R. C. Reddy, S. Akulapalli, V. J. Thannickal, T. J. Standiford, P. C. Andrews, and G. S. Omenn (2006): "Differential protein ex-

- pression profiling for iTRAQ-2DLC-MS/MS of lung cancer cells undergoing epithelial-mesenchymal transition reveals a migratory/invasive phenotype,” *Journal of Proteome Research*, 5, 1143–1154.
- Kuo, L. and B. Mallick (1998): “Variable selection for regression models,” *Sankhya: The Indian Journal of Statistics, Series B (1960-2002)*, 60, 65–81.
- Oberg, A. L., D. G. Mahoney, J. E. Eckel-Passow, C. J. Malone, R. D. Wolfinger, E. G. Hill, L. T. Cooper, O. K. Onuma, C. Spiro, T. M. Therneau, and B. H. R (2008): “Statistical analysis of relative labeled mass spectrometry data from complex samples using ANOVA,” *Journal of Proteome Research*, 7, 225–233.
- Oberg, A. L. and D. W. Mahoney (2012): “Statistical methods for quantitative mass spectrometry proteomic experiments with labeling,” *BMC Bioinformatics*, 13, S7.
- O’Hara, R. B. and M. J. Sillanpaa (2009): “A Review of Bayesian Variable Selection Methods: What, How and Which,” *Bayesian Analysis*, 4, 85–118.
- Perkins, D. N., D. J. C. Pappin, D. M. Creasy, and J. S. Cottrell (1999): “Probability-based protein identification by searching sequence databases using mass spectrometry data,” *Electrophoresis*, 20, 3551–3567.
- Plummer, M. (2003): “JAGS: A program for analysis of Bayesian graphical models using Gibbs sampling,” in *Proceedings of the 3rd International Workshop on Distributed Statistical Computing*, Vienna, Austria, <http://www-ice.iarc.fr/~martyn/software/jags/>.
- ProteoRed (2010): “ProteoRed Multi-centric Experiment 5,” http://www.proteored.org/PME5_main.asp.
- Ross, P. L., Y. N. Huang, J. N. Marchese, and B. Williamson (2004): “Multiplexed protein quantitation in *Saccharomyces cerevisiae* using amine-reactive isobaric tagging reagents,” *Molecular Cellular Proteomics*, 3, 1154–1169.
- Sidak, Z. (1968): “On multivariate normal probabilities of rectangles: their dependence on correlations,” *Ann. Math. Statist.*, 39, 1425–1434.
- Smyth, G. K. (2004): “Linear models and empirical Bayes methods for assessing differential expression in microarray experiments,” *Statistical Applications in Genetics and Molecular Biology*, 3:3.
- Thompson, A., J. Schalker, K. Kuhn, S. Kienle, J. Schwarz, S. Gunter, T. Neumann, and C. Hamon (2003): “Tandem mass tags: A novel quantification strategy for comparative analysis of complex protein mixtures by MS/MS,” *Analytical Chemistry*, 75, 1895–1904.
- Wang, X., G. A. Anderson, R. D. Smith, and A. R. Dabney (2012): “A hybrid approach to protein differential expression in mass spectrometry-based proteomics,” *Bioinformatics*, 28, 1586–1591.

Appendix

The JAGS model implementing the differential protein expression model described in section 2.2 is shown below.

```
model {
  Tau ~ dgamma(a.tau,b.tau)
  Sigma <- 1/sqrt(Tau)

  for (j in 1:P) {
    for (k in 1:m[j]) {
      Alpha[moffset[j]+k] ~ dnorm(a.alpha,b.alpha)
    }
  }

  for (e in 1:E) {
    kappa[e,1] <- 0
    for (s in 2:numberOfSamples[e]) {
      kappa[e,s] ~ dnorm(a.kappa,b.kappa)
    }
  }

  for (j in 1:P) {
    p[1,j] <- 0.0
    Beta[1,j] <- 0.0
    Gamma[1,j] <- 0.0
    for (g in 2:G) {
      p[g,j] ~ dbeta(a.p,b.p)
      Beta[g,j] ~ dbern(p[g,j])
      Gamma[g,j] ~ dnorm(a.gamma,b.gamma)
    }
  }

  for (n in 1:N[j]) {
    intensity[offset[j]+n] ~ dnorm(kappa[experiment[offset[j]+n],
      sample[offset[j]+n] - sampleoffset[experiment[offset[j]+n]]]
      + Alpha[peptide[offset[j]+n]]
      + Beta[group[offset[j]+n],j]*Gamma[group[offset[j]+n],j]
      ,Tau);
  }
}
```


Package Installation and Getting Started

The dpeaqms package is an R package which relies on the rjags package and the native JAGS library already being installed. The JAGS library can be downloaded and installed from <http://www-ice.iarc.fr/~martyn/software/jags/>. The rjags package is an R interface to this native JAGS library and can be installed within R by using the command

```
install.packages("rjags")
```

and then the dpeaqms package can be installed by using the command

```
install.packages("dpeaqms",  
                 repos="http://r-forge.r-project.org")
```

Once the dpeaqms package is installed, an overview vignette can be accessed by using the command

```
vignette("dpeaqms")
```

A vignette showing the analysis of the simulated dataset described in section 3 can be accessed using the command

```
vignette("dpeaqms.simulatedDataset")
```

The MCMC sampling scheme

The MCMC scheme is a Gibbs sampler which involves simulating realisations of the model parameters in turn from their full conditional distributions as follows:

- The log-normalisation ratio for reporter ion i in experiment e is simulated using $\kappa_{egi} | \cdot \sim N(A_{egi}, 1/B)$ for $(g, i) \neq (g_e, 1)$, where

$$A_{egi} = \frac{a_{\kappa} b_{\kappa} + \sum_{jk} (y_{egjki} - \alpha_{jk} - \beta_{gj} \gamma_{gj}) / \sigma^2}{b_{\kappa} + n_{\kappa} / \sigma^2}, \quad B_{ei} = b_{\kappa} + n_{\kappa} / \sigma^2,$$

and $n_{\kappa} = \sum_j m_j$ is the number of log-intensity measurements for reporter i in experiment e .

- The mean log-expression level for MS/MS spectrum k for (control) group 1 is simulated using $\alpha_{jk}|\cdot \sim N(C_{jk}, 1/D)$, where

$$C_{jk} = \frac{a_\alpha b_\alpha + \sum_{egi} (y_{egjki} - \kappa_{egi} - \beta_{gj} \gamma_{gj}) / \sigma^2}{b_\alpha + n_\alpha / \sigma^2}, \quad D = b_\alpha + n_\alpha / \sigma^2,$$

and $n_\alpha = \sum_{eg} n_{eg}$ is the number of log-intensity measurements for each mass spectrum k assigned to protein j .

- The parameter β_{gj} ($g \neq 1$) indicating whether or not the protein j is differentially expressed for group g with respect to control group 1 is simulated using probabilities

$$\begin{aligned} \pi(\beta_{gj} = 0|\cdot) &\propto (1 - p_{gj}) \exp \left\{ -\frac{1}{2\sigma^2} \sum_{eki} (y_{egjki} - \kappa_{egi} - \alpha_{jk})^2 \right\}, \\ \pi(\beta_{gj} = 1|\cdot) &\propto p_{gj} \exp \left\{ -\frac{1}{2\sigma^2} \sum_{eki} (y_{egjki} - \kappa_{egi} - \alpha_{jk} - \gamma_{gj})^2 \right\}. \end{aligned}$$

- The probability of differential expression of a protein j in group g is simulated using $p_{gj}|\cdot \sim \text{Beta}(a_p + \beta_{gj}, b_p + 1 - \beta_{gj})$ for $g \neq 1$.
- The difference in the mean log-expression levels of protein j between group g ($\neq 1$) and the control group is simulated using $\gamma_{gj}|\beta_{gj} = 0, \cdot \sim N(a_\gamma, 1/b_\gamma)$, its prior distribution, or $\gamma_{gj}|\beta_{gj} = 1, \cdot \sim N(E_{gj}, 1/F_{gj})$ as appropriate, where

$$E_{gj} = \frac{a_\gamma b_\gamma + \sum_{eki} (y_{egjki} - \kappa_{egi} - \alpha_{jk}) / \sigma^2}{b_\gamma + n_{gj}^\gamma / \sigma^2}, \quad F_{gj} = b_\gamma + n_{gj}^\gamma / \sigma^2$$

and $n_{gj}^\gamma = m_j \sum_e n_{eg}$ is the total number log-intensities measurements for samples in group g for protein j .

- The error standard deviation is simulated using $\sigma^{-2}|\cdot \sim \text{Ga}(a_\sigma + n/2, b_\sigma + G/2)$ where $G = \sum_{egjki} (y_{egjki} - \kappa_{egi} - \alpha_{jk} - \beta_{gj} \gamma_{gj})^2$ and $n = E n_I \sum_j m_j$ is the total number of log-intensity measurements for the isobaric labels.